DSTO

Topological Considerations of ATM Networks

T.A. Au

DSTO-TR-0799

19990607 033

Topological Considerations of ATM Networks

T.A. Au

Communications Division
Electronics and Surveillance Research Laboratory

DSTO-TR-0799

ABSTRACT

The key to achieving universal network scalability is hierarchical organisation, where routing information is aggregated between levels in the hierarchy. This approach can considerably reduce the overheads due to storage, transfer and processing of routing information for large networks. Other benefits include increased security and ease of management of the network. Due to the high routing complexity of the PNNI routing protocol, an optimal hierarchical structure is essential to reduce the huge overheads incurred. This report considers two hierarchical routing models based on the optimisation of the size of topology databases and routing complexity, respectively. The report also suggests some design guidelines for an hierarchical PNNI structure of a network topology.

RELEASE LIMITATION

Approved for public release

D E P A R T M E N T   O F   D E F E N C E
◆
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

Topological Considerations of ATM Networks

Executive Summary

In asynchronous transfer mode (ATM) networks, a dedicated virtual circuit is required to be set up from source to destination, which consists of a number of transmission links connected by intermediate switches. For the support of this connection-oriented service, signalling is essential to allow the transfer of service-related information in real time between the user and the network, and among the network entities. In particular, the process of selecting a suitable path substantially affects a number of performance measures associated with the existing and incoming calls.

Hierarchical network organisation is the key to achieving universal scalability. All nodes of a network are divided hierarchically into logical groups such that a node has complete routing information on its peer group but only abstract routing information on the rest of the network. This hierarchical routing approach can considerably reduce the overheads due to storage, communication and processing of routing information for large networks. Keeping the internal information of peer groups from others also results in increased security and ease of management of the network. Nevertheless, due to the lack of complete information about the network status, the hierarchical routing scheme may lead to a non-optimal routing decision (ie. increased path length), and hence performance degradation, in terms of longer end-to-end delay and reduced throughput.

Two hierarchical routing models are described in this report based on the optimisation of the size of topology databases and of the routing complexity, respectively. In support of the delivery of messages to their correct destination, the routing tables are usually pre-calculated for predetermined classes of quality of service (QoS). The complexity of these calculations is directly related to the size of the topology databases. It follows that a model for optimal topology databases is required to reduce this complexity. When an arriving call does not fall into any pre-calculated QoS class, a special on-demand routing computation is used whose complexity can be evaluated using the model for optimal computational complexity.

In particular, the routing complexity of private network-network interface (PNNI) implementations is likely to be higher than the original complexity for simple shortest paths, mainly due to the requirements for QoS. As the routing algorithm becomes more complex, the saving with an optimal hierarchical topology is even higher. Thus, an optimal PNNI hierarchical structure definitely results in huge complexity savings. To facilitate the migration of the current Defence Switched Data Network (DSDN) to a full-fledged ATM network, the report concludes with some design guidelines for a hierarchical PNNI structure of a network topology.

Authors

T. A. Au
Communications Division

Andrew Au is a Research Scientist in Network Architecture Group with interests in traffic control and performance guarantees for high speed networks. He has been investigating ATM signalling and communications issues in distributed computing environments.

Contents

1. Introduction

With the advent of optical fibre and fast switching technologies, future integrated broadband networks are expected to support a wide spectrum of services, among which asynchronous transfer mode (ATM) is selected to be the technology standard by the International Telecommunications Union Telecommunication Standardisation Sector (ITU-T). Using a connection-oriented packet switching technique, ATM combines both flexibility of packet switching and quality of service (QoS) of circuit switching. Similar to a circuit-switching network, an ATM network requires a dedicated virtual circuit to be set up from source to destination, consisting of a number of transmission links connected by intermediate switches. Once established, this virtual circuit is able to transmit all the data cells along the same path.

Signalling is an essential component of ATM networks for the support of connection-oriented services. In general, a communication network can be viewed as comprising of a signalling system and a transmission facility. The signalling system is responsible for call processing and associated functions like route selection, connection management, and QoS computation; whereas the transmission facility is responsible for the transport of information associated with the admitted calls.

Incoming call requests are first processed by the signalling system. In connection-oriented networks, path selection is performed only at call establishment phase. If a call is successfully set up, transmission resources will be allocated to the call throughout its connection time. These call requests are transmitted only on dedicated channels and do not contend for transmission media. They are not lost in the switches and their transmission delays are negligible due to the high transmission rate allocated for these signalling channels.

The process of selecting a suitable path substantially affects a number of performance measures of both the existing and new calls. Two types of performance measures are important in ATM networks: call level QoS, and cell level QoS. At the call level, performance objectives are usually measured in terms of blocking probability, and the time required setting up a call. This is the time for the routing algorithm to establish communication paths between source and destination across the network. At any time, transport capacity decides not only the number of calls admitted to the network, but also the cell level QoS of admitted calls. In particular, the routing algorithm works in conjunction with call admission control to ensure efficient use of network resources and guarantee a minimum level of QoS for the users, including cell loss, cell throughput, cell delay, and cell delay variation.

The ability of a network to admit calls is dominated by the capacity of the network, and the processing power of the signalling system, whose performance is crucial in determining the scalability of a network. The call set-up delay is mainly decided by the processing capacity of the signalling system where all set-up messages are first

processed. In some cases, call requests that are not processed immediately are allowed to wait or queue in the signalling system. Eventually, call requests are either admitted or rejected based on the availability of transmission resources. In general, the call blocking probability is decided both by the processing capacity of the signalling system and by the transmission capacity of the network.

When each virtual circuit is being set up, a routing decision is made to choose an end-to-end path between source and destination. All cells belonging to the virtual circuit subsequently traverse the ATM network following this path. These cells arrive at the destination node in the sequence in which they were transmitted. The signalling delay is one of the important performance measures at the call level. Even when transport capacity is available, a network can lose offered calls if the call set-up delay exceeds certain limits. Although provisioning permanent virtual path connections can reduce the call set-up delay by eliminating the processing time of signalling messages at some intermediate nodes, significant wastage of network capacity is inevitable.

In ATM networks, a guaranteed QoS is provided by means of a call level admission control mechanism, and the reservation of resources on a link-by-link basis. Thus, a call request is accepted if and only if, for each link on the route, adding the new connection will not violate the QoS guarantees of all existing connections sharing the same link. If a route is rejected, another will be selected if possible, between the same source and destination nodes trying to set up the call. Hence, the call routing and admission control policies can affect call set-up time and accepted call throughput significantly. Since these functions are highly computationally intensive, to assure accepted calls of a guaranteed QoS, the bottleneck for call set-up performance is likely to shift from the communication links to the processing elements. The processing requirement is even more demanding in defence networks where QoS is expected to meet specific military needs. Military value is often the norm of communication regardless of network load, necessitating the implementation of priority and pre-emption policies for resource management.

This report considers the design of a hierarchical routing structure in ATM networks. The benefits and potential drawbacks of such a network hierarchy are discussed in Section 2. Section 3 presents the PNNI protocol and the mechanisms used in distributing routing information. The routing overheads are considered in Section 4, where two hierarchical models are described based on the optimisation of the size of topology databases and of the routing complexity, respectively. Section 5 attempts to suggest some design guidelines for hierarchical clustering in ATM networks. Finally, Section 6 concludes the report.

2. Hierarchical Routing in ATM Networks

A routing protocol is required to distribute network control information through an ATM network, enabling each node to have sufficient topology information to process a call set-up request. Three major components are generally involved in a routing protocol: routing information, topology database and routing algorithm. The routing information is the information exchanged among network nodes so that each node can build up its own topological view of the network, in terms of the topology database at that node. The routing algorithm can then be run over this database to determine the best paths for different destinations. The best path to each reachable destination is pre-calculated for each service category, and kept in the form of a routing table[1] at each node. In support of the delivery of messages to their correct destination, a routing table is the table of current pre-calculated routes to other nodes associated with each service category. In the Private Network-Network Interface (PNNI) protocol, the routing table maintains a set of all Designated Transit List (DTL) stacks used for the pre-calculated routes. A DTL is a list of node and optionally link identifiers that completely specify a route to destination, which are inserted into signalling request messages.

The process of route determination involves an enormous routing overhead. The resources consumed include the amount of storage for the topology database, processing capacity to compute routing tables, and the bandwidth to propagate the topology database updates to other nodes. Since the topology database contains information about all resources in the network, the entire routing overhead increases with the size of the network.

2.1 Improved Scalability through Information Aggregation

The key to achieving universal scalability is hierarchical network organisation, with summarisation of reachability information between levels in the hierarchy. In simple ATM networks, establishing a hierarchy may not be necessary; however, a hierarchy becomes essential in larger networks. In flat routing, or non-hierarchical routing, network nodes have peer relationship with each other. Each node has to maintain the entire topology of the network, including resource availability for every physical link and reachability information for every node in the network. Therefore, the topology database at each switch contains complete information about the network, typically a graph with a vertex for every switch, an edge for every link, and a QoS attribute for every vertex and edge indicating resource availability (eg., available bandwidth). Each node must have a separate entry in its topology database for every node in the network. Thus the size of the topology database is proportional to the total number of

---

[1] The routing table is not the VPI/VCI swapping table. The latter is a table of VPI and VCI values related to each physical link that an ATM switch is attached. This is used for swapping the VPI and VCI fields of ATM cells passing through the switch, so that incoming cells are transmitted on the appropriate outbound links.

nodes in the network. Furthermore, a shortest path calculation is performed to configure the routing tables for the entire network.

As networks grow in size, the topology database at each node grows linearly with the number of nodes. Not only is memory consumed by increasing databases, but also more processing time is needed to determine routes. More bandwidth is needed to exchange routing information among the network nodes, which represents a significant burden on the communication lines. This characteristic of flat routing creates enormous overhead for large networks. Thus the main drawback of the non-hierarchical approach is excessive overhead for the storage, processing, and transfer of routing information. As the network grows to a certain size, it is no longer feasible for every node to have an entry for every other node, so the routing will have to be done hierarchically, as in telephone networks and the Internet.

Through aggregation of information, the hierarchical approach reduces the overheads of the routing algorithm for large networks. Nodes are grouped hierarchically into clusters, or logical groups. A peer group is a collection of nodes that all obtain the identical topology database and exchange full link state information with each other. While all nodes within a peer group have complete state information on each other, peer groups cannot be extended too widely since this would lead to excessive traffic and processing of routing information. Hence, peer groups are organised hierarchically and are associated with a higher-level parent peer group. All the details about how to route cells to destinations within its group is known, but the knowledge about the topological structure of other groups is limited. Therefore, each group maintains a view that contains its subgroups, the peer groups in its parent group, and so on. Each node is required to maintain a view of the network, which is periodically refreshed by routing updates. Such a view becomes extremely complex in the PNNI protocol because dynamic QoS attributes are involved, such as current load, and types of traffic in terms of service category. The ATM Forum has specified five service categories, including Constant Bit Rate (CBR), Real-Time Variable Bit Rate (rt-VBR), Non-Real-Time Variable Bit Rate (nrt-VBR), Unspecified Bit Rate (UBR), and Available Bit Rate (ABR).

2.2 Increased Security and Ease of Management

To reduce the amount of routing information distributed throughout the network, each higher level of the network hierarchy performs a complex process of summarising and compressing topology state information. Only summarised routing information is exchanged between related peer groups. Keeping the internal information of peer groups from others results in increased security of the network. In that case, each node is allowed to route within its own group according to the specific routing policy of that group. In particular, the use of hierarchical routing is essential in networks that are composed of subnetworks. By hiding the details of each subnetwork from the rest of the entire network, the management workload can be substantially reduced.
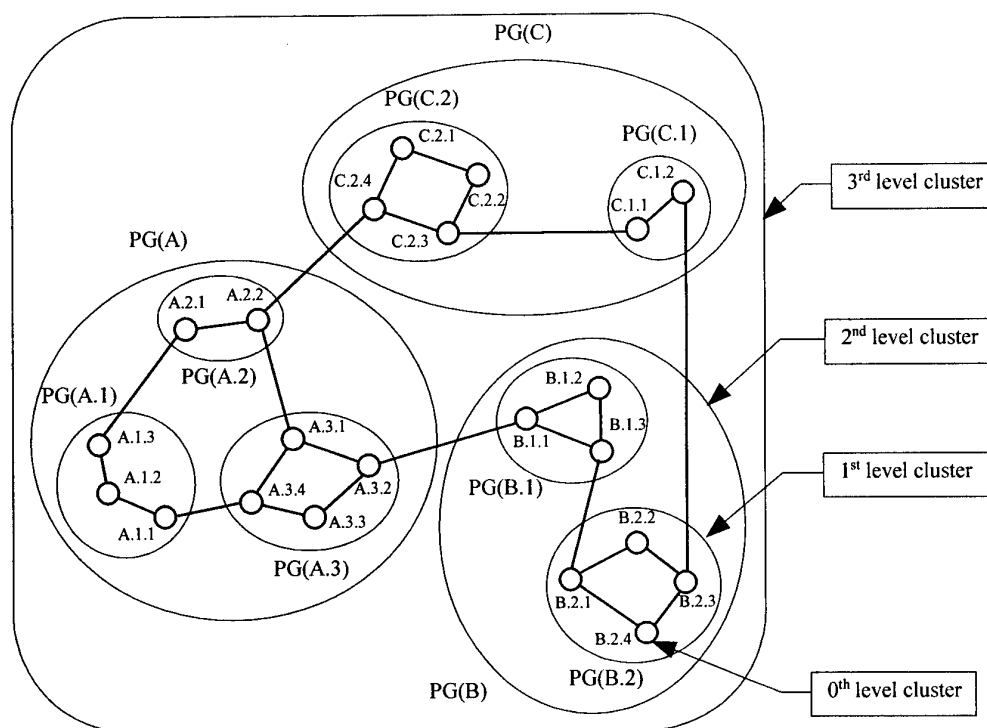
Figure 1: A 3-level clustered 24-node network

2.3 An Example

Essentially, the hierarchical routing procedure need not imply a hierarchical relationship between ATM switches. When applied to a flat network topology, this routing procedure provides very significant improvements as exemplified in the PNNI protocol. The network topology itself could certainly include a hierarchical structure as well [1]. Consider a network of a set of nodes (defined as 0^{th} level clusters), an m-level hierarchical clustering consists of grouping all the nodes into 1^{st} level clusters, which in turn are grouped into 2^{nd} level clusters and so on, until the m^{th} level cluster. The m^{th} level cluster is the highest level cluster that includes all the nodes of the network. An example is shown in Figure 1 in which a 3-level hierarchical clustering is imposed on a network of 24 nodes. This hierarchical classification scheme then leads to the tree representation as shown in Figure 2.

2.4 Some Drawbacks

In hierarchical routing, the savings in storage for topology database increases proportionally with the number of groups or levels. The routing overhead generated in hierarchical routing is much less than that in flat routing by maintaining less accurate views of the network. Due to the lack of complete information about the network status
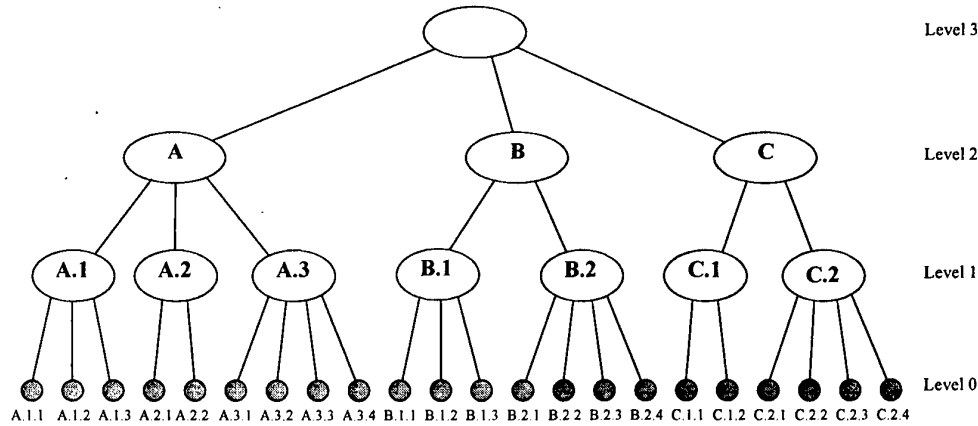
Figure 2: A tree representation of a 3-level clustered network

through topology aggregation, the hierarchical approach cannot accurately reflect the actual state of the network. This may lead to non-optimal routing decisions (ie, increased path length), and hence performance degradation. The average path length for hierarchical routing is more than that for flat routing, and longer path results in longer end-to-end delay. In addition, the use of extra resources affects both blocking and throughput performance, especially at high loads. Although allowing rerouting may slightly alleviate the call blocking performance, the call set-up delay is inevitably increased as a result of successive retry attempts.

At a higher hierarchical level, an aggregate QoS attribute of each group is obtained by applying an aggregation operation on all the QoS attributes of its members or subgroups. This aggregation operation contributes significantly to the complexity of selecting a path during a call request. With attributes optimistically aggregated, it is likely that the selected paths do not pass the admission control test due to insufficient resources during connection set-up. On the other hand, overly pessimistic aggregation hampers the selection of a suitable path with adequate resources even though such a path indeed exists [2].

3. The PNNI Protocol

The private network-network interface (PNNI) [3] is an ATM Forum specification for the protocols between switches in private ATM networks. The PNNI Phase 1 specification includes both routing and signalling functions. The PNNI routing protocol is used to reliably distribute network topology information so as to select a suitable path to any addressed destination. On the other hand, the PNNI signalling protocol is used for the connection establishment, management, and termination between ATM nodes within the network.

The PNNI Phase 1 protocol can be applied both to small networks of a few switches and to large global networks comprising millions of switches. The routing hierarchy is designed to reduce the routing overhead while providing efficient routing. In the PNNI protocol an entire ATM network is divided, at the lowest level of the hierarchy, into groups of smaller numbers of physical nodes (or switches) called peer groups. These groups are represented at higher levels as logical group nodes (LGNs) that may again be divided into parent groups if necessary, up to the highest level in the hierarchy. A LGN is an abstraction of a peer group for the purpose of representing that peer group in the next level of the PNNI routing hierarchy. The number of levels and the size of each group depend upon the original network size.

In the PNNI protocol, there is no direct exchange of routing information between two nodes that do not belong to the same peer group. Instead, each switch is required to maintain a view of the network status that is periodically refreshed by routing updates. In particular, this view is extremely complex since dynamic information is also included, such as current load and types of traffic on the links. A switch is required only to have very detailed topology and state information about switches in its immediate vicinity. Detailed planning is a cooperative process in which many switches along the route participate. This route is made more precise with added details, as the call request is forwarded and processed by the switches along the chosen route. Hence, the topology database in every physical or logical group node reflects only a partial view of the network, including the detailed topology information about the peer group in which the node resides, plus more abstract topology information representing the remainder of the PNNI routing domain.

3.1 ATM Addressing

The key to constructing the PNNI routing hierarchy is in the ATM addressing and identification based on the ATM end system addresses (AESAs). These addresses are positioned in a hierarchical relationship with one another. This differs from a flat relationship where the address of each switch does not bear any topological meaning. These AESAs are 20 octets long to support a large number of levels in the network hierarchy. Excluding the end system identifier[2] (ESI) and the selector[3] (SEL) fields, there is a maximum of 104 bits in the 13 high-order octets of the ATM address. At the extreme of 1 bit per level, the network can be administered by creating as many as 104 hierarchical levels. A logical ATM addressing structure provides a logical address path through the network, which makes ATM call routing simpler and more efficient. In practice, three to five hierarchical levels will most likely be used in large and even global networks. For example, a recommendation of two options for a defence ATM address plan has been proposed [4], which adopts three levels of topological hierarchy

---

[2] The end system identifier is used to identify a unique end system within a specified network to avoid address contention.
[3] The selector field is used to distinguish between different higher layer protocols and applications inside the end system.

for the current Defence Switched Data Network (DSDN). Two additional routing domains are reserved for future expansion, supporting up to five levels.

3.2 Information Exchange in PNNI

PNNI signalling is carried across network node interface (NNI) links on the same virtual channel, VCI=5, which is also used for signalling across the user network interface (UNI) to end systems. This signalling virtual channel either uses VPI=0, VCI=5 in case of non-associated signalling, or uses VCI=5 in the appropriate virtual path connection (VPC) in case of associated signalling. For the regular update of external reachability and link characteristics respectively, the PNNI Hello packets and PNNI topology state packets (PTSPs) are exchanged at intervals on a timed basis over PNNI routing control channels (RCCs) between switches that are logically or physically adjacent. These are reserved virtual channel connections (VCCs) with VPI=0 and VCI=18 for physical links, or VCI=18 within the appropriate VPI value for logical links. For the exchange of PNNI routing protocol packets between LGNs at higher levels, a switched VCC (SVCC) is established through unrelated nodes at lower levels whose VPI and VCI values are assigned by signalling in a normal manner. Mechanisms such as flooding, sequence numbers, lock-step acknowledgments, and checksums are used to ensure reliable and timely delivery of PTSPs.

PNNI Hello packets are used to discover and verify the identity of neighbour nodes and to determine the status of the links to those nodes. The PNNI Hello protocol supports the exchange of peer group identifiers (PGIDs) to determine whether the neighbours belong to its peer group or not. Each node collects its local state information in PNNI topology state elements (PTSEs), including the identity and peer group membership of the node's immediate neighbours, and the status of its links to the neighbours. PTSEs are encapsulated within PTSPs, which are reliably flooded among all LGNs within a peer group. PTSPs are sent at regular intervals or when triggered by a significant event such as a major change in any topology information group. The topology database of a node consists of a collection of all PTSEs received, representing the node's present view of the PNNI routing domain.

3.3 Route Generation

The PNNI protocol is based on source routing, in which a path is computed in the source node across the ATM network to a destination node. This route is represented by a list of nodes on the path called a Designated Transit List (DTL), which is included in the signalling messages to establish connections. A node along the path from the source to the destination may not have sufficient resources to support the requested connection. This is possible if the topology information used to compute a path is not up to date due to the delay incurred in the periodic update of topology status. In such cases, a crankback mechanism is used to clear the call back to the node that created and inserted the DTL into the connection request. These include the original source node, and the ingress border nodes of each peer group along the path. If an intermediate

node rejects the call, the call is rerouted back along the path that it followed to such a node to insert a DTL. If possible, this node then recalculates a new path across its own group, avoiding the node that rejected the call, and reforwards the request. These nodes maintain state information about all requests that they have forwarded until the connection is established, or a connection reject message is received from the destination end system.

3.4 Pre-calculated Routing versus On-demand Routing

For large networks, the hierarchical routing model can substantially reduce the overheads due to storage, transfer and processing of routing information. Due to the complexity of the PNNI protocol, on-demand route calculation can take a considerable time, thus increasing the latency of connection establishment. Alternatively, pre-calculated paths can be computed periodically to generate reasonable routes for a given service category based on some optimisation criteria, usually after each periodic update of topology information. The process of route generation can be run in the background to determine routes to all possible destinations. It produces a set of pre-calculated routes in terms of DTLs for locally originated calls to the destinations reachable at the given node. When the set of pre-calculated paths does not meet the QoS parameters of a given call, it is necessary to attempt on-demand calculation of a route. The latter is suitable either for originating a call, or for propagating a call with a given DTL.

It is important that an optimal route is located for virtual connections with high bandwidth-duration product (ie, bandwidth × duration) so as to minimise the allocated network resources. For virtual connections with low bandwidth-duration product, minimising the route determination processing becomes more important than finding an optimal route, as it is expected that such connections may constitute a large portion of the total usage of ATM networks.

4. Routing Overheads in ATM Networks

On-demand routing produces more computation overhead than pre-calculated routing, and this is tolerable at low loads. At higher loads, this overhead becomes significant, but pre-calculated routing may require greater rerouting because the routing tables cannot always reflect all the possible paths at a particular moment. In the case of on-demand routing, the computation overhead is related to the inter-arrival time of incoming calls. While in the case of pre-calculated routing, the computation overhead is associated with the update period of synchronising PNNI topology information among network nodes. Selecting the update period according to the call inter-arrival time ensures that topology information at every node is relatively recent for call processing. In general, on-demand route computation can reduce the call set-up delay by pruning the links that will most likely reject the call from the topology,

resulting in better performance than pre-calculated route computation. Thus, the use of on-demand route computation mechanism is preferred.

Due to the greater variability of the QoS-based routing metrics, the computation overhead in the PNNI protocol is large. Using commercial processors for PNNI calculation will likely support very limited rates of call set-up requests. In addition, each connection set-up request could experience significant call set-up latencies, perhaps exceeding hundreds of milliseconds within large networks, not to mention the prolonged propagation delay across satellite links. These set-up latencies could significantly degrade perceived network responsiveness.

Through a hierarchical partitioning of a network, the PNNI protocol is designed to reduce the overhead of maintaining the entire topology at each node in order to provide for efficient routing. Essentially, the hierarchical topology of a network should be carefully structured in order to achieve minimal routing complexity[4]. Typically, the calculation of shortest paths in routing always involves high computational complexity. For instance, Dijkstra's algorithm[5] results in an $O(N^2)$ computational complexity, where N is the total number of nodes in the network.

4.1 Storage and Communications Overheads of Routing Information

Hierarchical routing can be applied to lower the routing cost through the reduction of storage for the topology database. Smaller topology databases leads to less storage and processing in the nodes and limits the inherent routing communications. Given a network with N nodes, the flat routing scheme requires each node to keep a topology database that is composed of N entries. The size of the topology database thus grows linearly with the number of nodes. Since the database size is directly related to the routing cost in terms of the communication capacity and the storage, it is important to determine the optimal hierarchical clustering structures that lead to a minimal database size.

Kleinrock and Kamoun [1] determine the optimal hierarchical clustering structures of the network nodes for large networks by minimising the size of the required topology databases. The degree of a k^{th} level cluster is defined as the number of k-1^{st} level clusters that the k^{th} level cluster encloses. The optimum solution has all clusters at all levels composed of the same number of lower level clusters. If n_k is the degree of all the k^{th} level clusters, then $n_k = N^{1/m}$, where m is the fixed number of levels in the hierarchy. Using a random geometric approach, Rougier and Kofman [5] also confirms the

---

[4] Complexity is the level of difficulty in solving mathematically posed problems as measured by the time, number of steps or arithmetic operations, or memory space required (called time complexity, computational complexity, and space complexity, respectively). The interesting aspect is usually how complexity scales with the size of the input (called the scalability), where the size of the input is described by some number N. Thus an algorithm may have computational complexity $O(N^2)$ in the order notation, in which case if the input doubles in size, the computation will take four times as many steps. The ideal is a constant time algorithm ($O(1)$) or failing that, $O(N)$.

[5] Dijkstra's algorithm is a common least-cost routing algorithm in use in packet-switched networks.

existence of this optimum in the case of a random network where peer groups are of variable sizes.

The minimum database size, l, is $mN^{1/m}$. The global optimal clustering is achieved when the number of levels is $m^* = \ln N$, and when all clusters at all levels have equal number of lower level clusters. It follows that the optimal cluster size is equal to the base of natural logarithm ($e \approx 2.718$). The corresponding minimum database size is $e \ln N$. Hence, enormous gains in the routing cost can be obtained whereby the size of the topology database may be reduced from N entries to the order of $e \ln N$ entries. For example, $e \ln N = 15.5$ for $N = 300$.

In general, the reduction of routing information inevitably leads to an increase in network path length. Fortunately, this increase vanishes in the limit of very large networks with vast number of nodes [1]. In other words, hierarchical routing for large networks will approach similar throughput-delay performance as the non-clustered routing, with relatively no significant increase in path length.

Although the application of the hierarchical routing scheme appears to result in a degradation of the performance of the network, the database reduction provides savings in capacity, storage, throughput and delay which more than compensate for the vanishing increase in path length. Further, most of the database reduction can be obtained with hierarchical clustering whose number of levels is quite a bit smaller than m^*, and the cost at a small m is quite minimal [1]. Hierarchical routing schemes operating with a small number of levels $2 \le m \le 4$ already yield substantial database reduction for a relatively small increase in path length.

4.2 Routing Complexity

Given a network of N nodes, Van Mieghem [6] investigates the structure of a PNNI topology with m hierarchical levels based on the criterion of routing complexity. An optimum solution is found to exhibit symmetry such that on each hierarchical level, the peer groups have the same number of nodes. In the optimal hierarchical topology, the degree of peer groups at each level increases with the hierarchical level k. This characteristic is also confirmed by Rougier and Kofman [5] using a random geometric approach. Importantly, the peer group size is the smallest at the physical or lowest level, which substantially simplifies the analysis of QoS and traffic issues.

Assuming a routing complexity of $O(N^2)$ that corresponds to Dijkstra's algorithm for shortest paths, the number of nodes within each group at level k, n_k, can be estimated in terms of the total number of nodes of the network, N, and the number of levels, m [5]:

$$n_k \approx N^{\frac{\left(\frac{4}{3}\right)^k}{3\left[\left(\frac{4}{3}\right)^m - 1\right]}}$$

In brief, the peer group sizes are related by $n_{k+1} \approx n_k^{4/3}$, for hierarchical level $k = 1, 2, \ldots,$ m-1. With this optimal topology, the complexity is indeed reduced at a higher layer of

11

hierarchical levels. However, even for very large networks with millions of nodes, the complexity does not evolve significantly for $m > 4$. Thus, a small number of hierarchical levels is sufficient to significantly reduce the complexity of the routing computations [5, 6].

The more complex the routing algorithm is, the higher is the saving with an optimal hierarchical topology. Compared to the original complexity of $O(N^2)$ for simple shortest paths, the routing complexity of PNNI implementations is likely to be higher than that of Dijkstra's algorithm, mainly due to the requirements for QoS. In that case, the optimal PNNI hierarchical structure definitely results in huge complexity savings [6].

5. Some Guidelines for Hierarchical Clustering

In this section, we attempt to put some important factors into perspective in relation to the design of an efficient hierarchical structure for an ATM network. Examples are the choice of the number of hierarchical levels, and the size of each group of nodes at each level.

5.1 Hierarchical Routing Model

In ATM networks, the routing tables are usually pre-calculated for predetermined classes of QoS. The complexity of these calculations is directly related to the size of the topology databases, thus the model derived by Kleinrock and Kamoun [1] is appropriate. When an arriving call does not fall into any pre-calculated QoS class, a special on-demand routing computation is performed. The routing complexity can be evaluated using the model derived by Rougier and Kofman [5]. Therefore, it is necessary to evaluate the connection mechanisms in relation to the ratio between the number of connection establishment using on-demand computations or pre-calculated routing tables.

Besides, it is extremely helpful to examine the characteristics of the majority of connections expected in an ATM network. On-demand calculation of an optimal route is beneficial to virtual connections with high bandwidth-duration product, whereas pre-calculated routing reduces processing of route determination for those with low bandwidth-duration product.

Nevertheless, the difference between these two models is not significant. Also, it is not possible in practice to implement too many hierarchical levels with the exact number of peer groups and members. On these grounds, 2 to 4 is the recommended number for levels in a hierarchy. As a rule of thumb, only a very small number of hierarchical levels, $m = O(\log_2 \log_2 N)$, are needed in realistic configurations [6]. For example, $\log_2 \log_2 N = 3.04$ for $N = 300$.

Table 1: The size of peer groups n_k at level k against the number of hierarchical levels m in the DSDN based on a minimum routing complexity

k \ m	2	3	4
4			8.1
3		11.8	4.8
2	26.0	6.4	3.2
1	11.5	4.0	2.4

The size of a peer group is limited by the memory and processing power of the switches in that peer group. As the size of the peer group grows, the amount of topology information to be processed increases, so does the processing required to compute paths in that network. In addition, the use of on-demand route computation mechanism in QoS-guaranteed networks becomes a great burden on processing power, it is recommended that the model based on routing complexity should be considered.

We take the current DSDN as an example where the number of switches is assumed to be $N = 300$. The optimal network topology according to the Rougier and Kofman model is given in Table 1, where the number of hierarchical levels m varies from 2 to 4, and n_k is the number of nodes in a peer group at level k. In fact, this topology can be generously supported by the two proposals for the defence ATM address plan [4].

5.2 Clustering Network Nodes

In the design of a clustered network structure, locality should be taken into special consideration by grouping together nodes that are relatively adjacent to each other at a lower hierarchical level. Distant nodes should only be linked at a later stage to avoid lengthy propagation delay in the periodic exchange of routing information. To avoid possible link-state staleness, routing information is expected to reach other nodes as soon as it is disseminated. Since the routing update period is normally longer at a higher level, clustering based on locality ensures that the time to propagate routing information is always negligible compared to the update period at any level [2].

At any hierarchical level, at least one path should exist between any pair of nodes within a peer group. Also, each group should completely contain the shortest path so that the traffic always follows paths internal to the group [1]. In other words, the shortest path between any two members of a group should not pass through any switches that are non-members. In so doing, calls local to a group do not need to route through a different group. Better performance of the hierarchical model is expected due to better connectivity of peer groups. Thus, the connectivity inside poorly connected groups should be increased accordingly to enhance the performance.
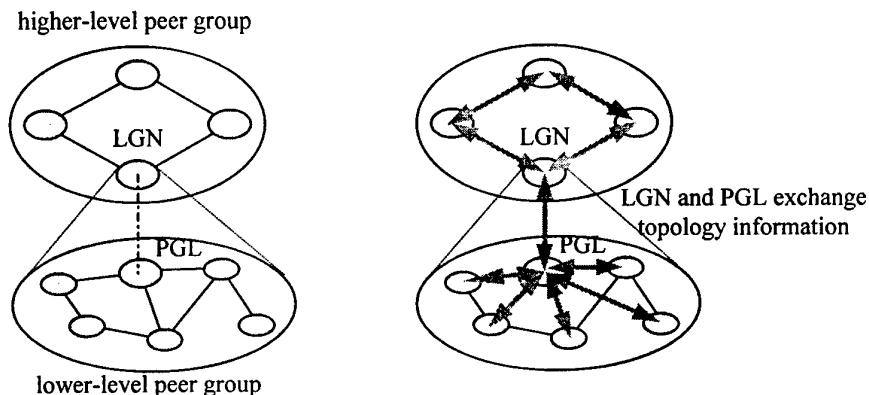
higher-level peer group

lower-level peer group

Figure 3: Peer group leader (PGL) and logical group node (LGN)

Besides, hierarchical routing is suitable for networks with minimal connectivity between isolated groups. It minimises routing overheads without affecting the performance. While for highly connected networks, the abstraction of routing information hides substantial amount of connectivity information and thus leads to degradation of performance through non-optimal routing solutions.

5.3 Selecting Peer Group Leader

The peer group leader (PGL) is a single switch that operates in multiple peer groups, its role is to convey aggregated routing information between the parent and child peer groups. In addition to being a normal member and an elected leader of its peer group, the PGL operates as a normal member of the parent peer group. Within a higher-level peer group, the peer group that a PGL represents is called a logical group node (LGN). This single LGN aggregates information about its child peer group to form a logical image of the peer group at higher levels. The duality of PGL and LGN is shown in Figure 3, indicating the flow of information [7].

Inevitably, the PGL is required to have significant memory and processing power to perform the aggregation and to participate in activities at multiple hierarchical levels. Apart from these operations, a PGL should have enough processing capacity to handle its own call requests and those forwarded via this node, which involves the highly computationally intensive functions of call routing and admission control.

5.4 Configuration of Virtual Path Topology

Given an ATM network, a set of virtual paths can be established across the physical network to create a virtual topology. This embedded topology is called the virtual path topology, or virtual path network. A virtual path is a logical direct link between two nodes that define a prearranged route possibly across several physical hops. The virtual path concept increases the logical connectivity of the ATM network, thus enhancing the performance of the physical hierarchy.

Instead of setting up end-to-end virtual circuit connections on a switch-by-switch basis, a virtual path network simplifies network management and control by reducing nodal processing demands at intermediate nodes. By considering only one-hop or two-hop virtual paths, the mechanism of path selection becomes very simple and fast while the routing tables are small. To reduce set-up delay, end-to-end virtual path connections are provisioned a priori in the network to carry a certain number of virtual circuits. Dedicated resources can be assigned to a virtual path for each service class between any pair of source and destination, within which all the virtual circuits are transported, processed and managed in a similar manner.

In some cases, the virtual path topology may be dynamically configured according to current network requirements. For instance, traffic characteristics may vary over different time frames (hour to hour, day to day, or month to month). This requires a flexible means of adding and deleting virtual paths to adjust to the variability of traffic over the different time frames that need to be provided in this environment.

Although the use of provisioned virtual paths reduces call set-up delay, this approach can result in significant wastage of network capacity due to a decrease in statistical multiplexing gains. Therefore, the problem of virtual path distribution should be addressed with extreme care, so as to tune the fundamental trade-offs between the call level QoS and the overall network performance [8].

6. Conclusions

Hierarchical network organisation is fundamental in supporting universal scalability of ATM networks. A hierarchical partitioning of a network entails grouping the nodes into clusters or peer groups so that reachability information is aggregated between levels in the hierarchy. This report has discussed this hierarchical routing approach and its impact on the performance of the network. Some design guidelines have also been suggested for a hierarchical PNNI structuring of a network topology.

For large networks, the hierarchical routing model can substantially reduce the overheads due to storage, transfer and processing of routing information. The routing traffic in a hierarchical network is considerably less than that in flat routing, resulting in increased network efficiency. Other benefits include increased security, and ease of management of the network. Nevertheless, due to the lack of complete information about the network status, the hierarchical routing approach may lead to non-optimal routing decisions (ie. increased path length), and hence performance degradation in terms of longer end-to-end delay and reduced throughput.

We have described two hierarchical routing models based on the optimisation of storage overhead and routing complexity, respectively. In ATM networks, the routing

tables are usually pre-calculated for predetermined classes of QoS. The complexity of these calculations is directly linked to the size of the topology database. When an arriving call does not fall into any pre-calculated QoS class, a special on-demand routing computation is used. By minimising this computational complexity, the optimum solution implies that the size of peer groups at higher level must be greater than that at lower level.

Before an optimal hierarchical structure is selected, it is necessary to evaluate precisely the ratio between the number of connection set-up requests using on-demand computations or pre-calculated routing tables. In real implementations, the routing complexity of the PNNI protocol is likely higher than that of Dijkstra's algorithm for simple shortest paths, mainly due to the requirements of quality of service. In that case, the PNNI hierarchical structure based on optimised routing complexity offers more advantageous complexity savings.

References

[1] L. Kleinrock, and F. Kamoun, "Hierarchical Routing for Large Networks: Performance Evaluation and Optimization," Computer Networks, vol. 1, no. 3, pp. 155-174, 1977.

[2] S. Ahn, G. Chittiappa, A. U. Shankar, "Numerical Evaluation of Hierarchical QoS Routing," Technical Report CS-TR-3905, University of Maryland, College Park, 1998.

[3] The ATM Forum, "Private Network-Network Interface Specification Version 1.0," March 1996.

[4] T. A. Au, and W. D. Blair, "Issues in Asynchronous Transfer Mode Network Addressing for the Australian Department of Defence," DSTO Technical Report DSTO-TR-0595, April 1998.

[5] J. L. Rougier, and Daniel Kofman, "Optimization of Hierarchical Routing Protocols," IFIP ATM Workshop, pp. 43/1-43/10, 1998.

[6] Piet Van Mieghem, "Estimation of an Optimal PNNI Topology," Proc. IEEE ATM '97 Workshop, pp. 570-7, 1997.

[7] H.J.R. Dutton, and P. Lenhard, "Asynchronous Transfer Mode (ATM) Technical Overview," Second Edition, Prentice Hall, 1995.

[8] Nikos G. Aneroussis, and Aurel A. Lazar, "Virtual Path Control for ATM Networks with Call Level Quality of Service Guarantees," IEEE/ACM Transactions on Networking, vol. 6, no. 2, pp. 222-36, April 1998.

Topological Considerations of ATM Networks

T. A. Au

AUSTRALIA

DEFENCE ORGANISATION

DGC3ID

S&T Program
    Chief Defence Scientist ⎫
    FAS Science Policy ⎬ shared copy
    AS Science Corporate Management ⎭
    Director General Science Policy Development
    Counsellor Defence Science, London (Doc Data Sheet)
    Counsellor Defence Science, Washington (Doc Data Sheet)
    Scientific Adviser to MRDC Thailand (Doc Data Sheet)
    Scientific Adviser Policy and Command
    Navy Scientific Adviser (Doc Data Sheet and distribution list only)
    Scientific Adviser - Army (Doc Data Sheet and distribution list only)
    Air Force Scientific Adviser
    Director Trials

    Aeronautical and Maritime Research Laboratory
    Director

    Electronics and Surveillance Research Laboratory
    Director (1 copy)
    Chief of Communications Division
    Research Leader Military Information Networks
    Head Network Architecture Group
    Author: T. A. Au

    DSTO Library
    Library Fishermens Bend
    Library Maribyrnong
    Library Salisbury (2 copies)
    Australian Archives
    Library, MOD, Pyrmont (Doc Data sheet only)

Capability Development Division
    Director General Maritime Development (Doc Data Sheet only)
    Director General Land Development (Doc Data Sheet only)
    Director General Aerospace Development (Doc Data Sheet only)

Army
    ABCA Office, G-1-34, Russell Offices, Canberra (4 copies)

SO (Science), DJFHQ(L), MILPO Enoggera, Queensland 4051 (Doc Data Sheet only)
NAPOC QWG Engineer NBCD c/- DENGRS-A, HQ Engineer Centre Liverpool Military Area, NSW 2174 (Doc Data Sheet only)

Intelligence Program
DGSTA Defence Intelligence Organisation

Acquisitions Program
PD DEFNET
PD DSDN

Corporate Support Program (libraries)
OIC TRS, Defence Regional Library, Canberra
*US Defence Technical Information Center, 2 copies
*UK Defence Research Information Centre, 2 copies
*Canada Defence Scientific Information Service, 1 copy
*NZ Defence Information Centre, 1 copy
National Library of Australia, 1 copy

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
 Library
 Head of Aerospace and Mechanical Engineering
Serials Section (M list), Deakin University Library, Geelong
Senior Librarian, Hargrave Library, Monash University
Librarian, Flinders University

OTHER ORGANISATIONS

NASA (Canberra)
AGPS
State Library of South Australia
Parliamentary Library, South Australia

OUTSIDE AUSTRALIA

ABSTRACTING AND INFORMATION ORGANISATIONS
Library, Chemical Abstracts Reference Service
Engineering Societies Library, US
Materials Information, Cambridge Scientific Abstracts, US
Documents Librarian, The Center for Research Libraries, US

INFORMATION EXCHANGE AGREEMENT PARTNERS
Acquisitions Unit, Science Reference and Information Service, UK
Library - Exchange Desk, National Institute of Standards and Technology, US

SPARES (5 copies)

Total number of copies: 64

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA		1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)

2. TITLE Topological Considerations of ATM Networks	3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)

4. AUTHOR(S) T.A. Au	5. CORPORATE AUTHOR Electronics and Surveillance Research Laboratory PO Box 1500 Salisbury SA 5108 Australia

6a. DSTO NUMBER DSTO-TR-0799	6b. AR NUMBER AR-010-881	6c. TYPE OF REPORT Technical Report	7. DOCUMENT DATE March 1999

8. FILE NUMBER E8709-4-19 PT 2	9. TASK NUMBER ADF96/295	10. TASK SPONSOR DGC3ID	11. NO. OF PAGES 16	12. NO. OF REFERENCES 8

13. DOWNGRADING/DELIMITING INSTRUCTIONS	14. RELEASE AUTHORITY Chief, Communications Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT *Approved for public release*   OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DIS NETWORK OFFICE, DEPT OF DEFENCE, CAMPBELL PARK OFFICES, CANBERRA ACT 2600

16. DELIBERATE ANNOUNCEMENT No Limitations

17. CASUAL ANNOUNCEMENT	Yes

18. DEFTEST DESCRIPTORS Asynchronous transfer mode, Computer networks, Hierarchies, Routing

19. ABSTRACT

The key to achieving universal network scalability is hierarchical organisation, where routing information is aggregated between levels in the hierarchy. This approach can considerably reduce the overheads due to storage, transfer and processing of routing information for large networks. Other benefits include increased security and ease of management of the network. Due to the high routing complexity of the PNNI routing protocol, an optimal hierarchical structure is essential to reduce the huge overheads incurred. This report considers two hierarchical routing models based on the optimisation of the size of topology databases and routing complexity, respectively. The report also suggests some design guidelines for an hierarchical PNNI structure of a network topology.